

Fake News Detection System Using MultinomialNB and Django Framework

Geeta Makhija^{1*}, Bhakti Gotarne², Rutuja Mane³, Bhakti More⁴

Abstract

The emergence of the World Wide Web and the rapid growth of online platforms have transformed the landscape of news dissemination. However, the rise of social media has also led to an overwhelming influx of potentially unreliable information, making it increasingly challenging to verify the truthfulness of articles. This verification process has become a daunting task, necessitating a thorough examination of various domain-specific aspects to ascertain the credibility of news content. In response to this issue, machine learning algorithms have shown considerable promise in automatically detecting fake news. Researchers are actively employing a diverse range of performance metrics to evaluate the efficacy of these algorithms. Natural language processing (NLP) techniques are essential in this process, as they aid in data pre-processing, thereby improving the accuracy of machine learning models. By leveraging extracted textual properties, researchers can train and evaluate machine learning classifiers designed to distinguish between genuine and fabricated content. This includes leveraging various features and performance metrics to evaluate the effectiveness of these classifiers. Ultimately, this study aims to contribute to the ongoing efforts to combat misinformation in the digital age by providing insights into the intersection of machine learning, NLP, and news verification.

Keywords: Machine learning, fake news, MultinomialNB, social media, term frequency-inverse document frequency (TF-IDF)

INTRODUCTION

Originally devised to augment user satisfaction, has recently emerged as a significant conduit for the propagation of misinformation [1]. This system persistently presents content akin to users' prior interactions, irrespective of its veracity, thereby perpetuating the circulation of false information [2]. Essentially, upon encountering misinformation, users are consistently steered towards similar content, with the system potentially fine-tuning the display of objectionable material in their social media feeds [3].

Previous studies have predominantly concentrated on identifying fake news by scrutinizing linguistic or structural cues such as authorial clarity or article length [4–7]. However, this method neglects to account for the preferences of users who consume or share fake news, as well as the intricate dynamics of social media networks, where misinformation proliferates. With the advent of sophisticated language models such as ChatGPT, differentiation between genuine and fake news using conventional methods has become increasingly arduous [8]. Consequently, there is an urgent need to reassess approaches to fake news detection.

*Author for Correspondence

Geeta Makhija
E-mail: geeta.lmakhija@gmail.com

¹Assistant Professor, Department of Computer Engineering, Parvatibai Genaba Moze College of Engineering, Wagholi, Pune, Maharashtra, India

²⁻⁴Student, Department of Computer Engineering, Parvatibai Genaba Moze College of Engineering, Wagholi, Pune, Maharashtra, India

Received Date: September 25, 2024
Accepted Date: December 23, 2024
Published Date: December 12, 2024

Citation: Geeta Makhija, Bhakti Gotarne, Rutuja Mane, Bhakti More. Fake News Detection System Using MultinomialNB and Django Framework. *Current Trends in Information Technology*. 2025; 15(1): 23–32p.

In this study, we propose an innovative detection framework that not only considers the visual attributes of content but also incorporates the

behaviors of users who propagate and disseminate fake news, along with the characteristics of the networks that facilitate its propagation. Moreover, advancements in AI technology have facilitated the creation of fake news that closely mimics authentic content, rendering it challenging to discern manipulated information based solely on surface features [9]. Hence, a more holistic approach encompassing user traits and social network dynamics is imperative to combat the proliferation of misinformation [10, 11].

As the popularity of fake news detection systems has increased, there has been a concurrent rise in the development of models designed to circumvent algorithmic identification methods employed by established fake news detection systems. This trend poses a significant challenge, as individuals intentionally craft misinformation that evades detection by these systems [12]. Adversarial attacks targeting fake news detection systems can impede the development of accurate and reliable detection mechanisms. The evolving technology utilized to disseminate fake news further complicates this issue, yet current detection systems primarily rely on the linguistic and compositional attributes characteristic of fake news, as in previous iterations. This ongoing arms race underscores the need for continuous innovation in the field of fake news detection to stay ahead of sophisticated tactics employed by purveyors of misinformation [13].

Our research endeavors to enhance the efficacy of fake news detection by addressing the shortcomings of prior investigations, which often disregard the characteristics of information recipients. We developed a comprehensive detection framework that amalgamates diverse content features, user behaviors on social media, and network structures involved in the dissemination of fake news.

LITERATURE SURVEY

Textual Content-Based

Previous research on news authentication has predominantly focused on analyzing textual elements and user metadata. Extracting statistical features from message content has been extensively discussed in the literature for identifying fake news [14]. These features include the recognition of distinct writing styles, emotional cues commonly associated with fake news, and factors such as network connections, style analysis, and individual emotions, all of which have shown promise in fake news detection. Researchers have explored how writing style affects readers' viewpoints and attitudes after reading such posts. Emotion, a significant predictor in many fake news detection studies, is often conveyed through user positions or simple statistical emotional features [15]. For example, one method introduced a new approach based on dual emotions to identify fake news, drawing insights from both publishers' and users' content along with emotional representation. Another study utilized a machine learning model employing convolution filters to differentiate between various levels of text information granularity for fake news identification, achieving an accuracy of 91% using an SVM model [16, 17].

Social Context Based

User interactions with news stories on social media can provide valuable supplementary insights beyond the content itself [18]. One approach is to use a knowledge graph to detect fake news based on the actual content. Another study utilized a graph kernel-based approach to uncover propagation patterns and attitudes. However, gathering social context features can be challenging because of their noisy and unstructured nature and the time-intensive process involved.

Stance Detection Overview

Stance detection involves determining the author's or text's perspective regarding a specified target, such as a topic, headline, or individual [19]. This process typically involves three factors and a machine learning-based categorization method to establish a comparison approach. Group titles, such as help, against, for, or neutral, are determined by the issue at hand and encompass various categories such as political debates, articles, and even internal corporate discussions. Opinion mining, particularly in detecting tweets or short-text stances, has garnered significant attention. Numerous languages have made collaborative efforts to provide datasets and foster research in this domain [20].

Misleading Headlines

Detecting misleading headlines involves categorizing each article’s treatment of the headline assertion into one of four categories: agree, discuss, disagree, or irrelevant (when the headline addresses a topic different from the body text). The emergence of annotated corpora and increased adoption of new technologies to combat the fake news challenge have introduced a fresh hurdle in fake news analysis [21, 22].

PROPOSED METHODOLOGY

This study employs a three-stage methodology. The initial stage encompassed pre-processing, in which data filtering and cleaning techniques were applied to extract semantic features from the raw dataset. The techniques included a stopwords filter to categorize the data by removing prepositions. In addition, HTML tags and non-English characters were removed to filter the data and eliminate impurities irrelevant to classification. In the second stage, semantic features are converted into feature vectors [23]. The final stage involved the utilization of machine learning and deep learning classifiers to group items in the dataset. Each method was separately applied to the same dataset, and the details of the three stages.

Pre-processing Stage

Data pre-processing, akin to text mining, involves converting unstructured textual data into structured data that is suitable for machine processing. The stopwords technique, which is a common method for data filtering, information retrieval, and text classification, was employed to clean the classified dataset. Stopwords, such as “the,” “in,” “a,” “an,” and “with,” among others, were removed [24, 25]. The Python Standard Library was utilized to remove HTML tags using the `remove_tags` function, followed by pre-processing of the text function to eliminate non-English characters, as shown in Figure 1.

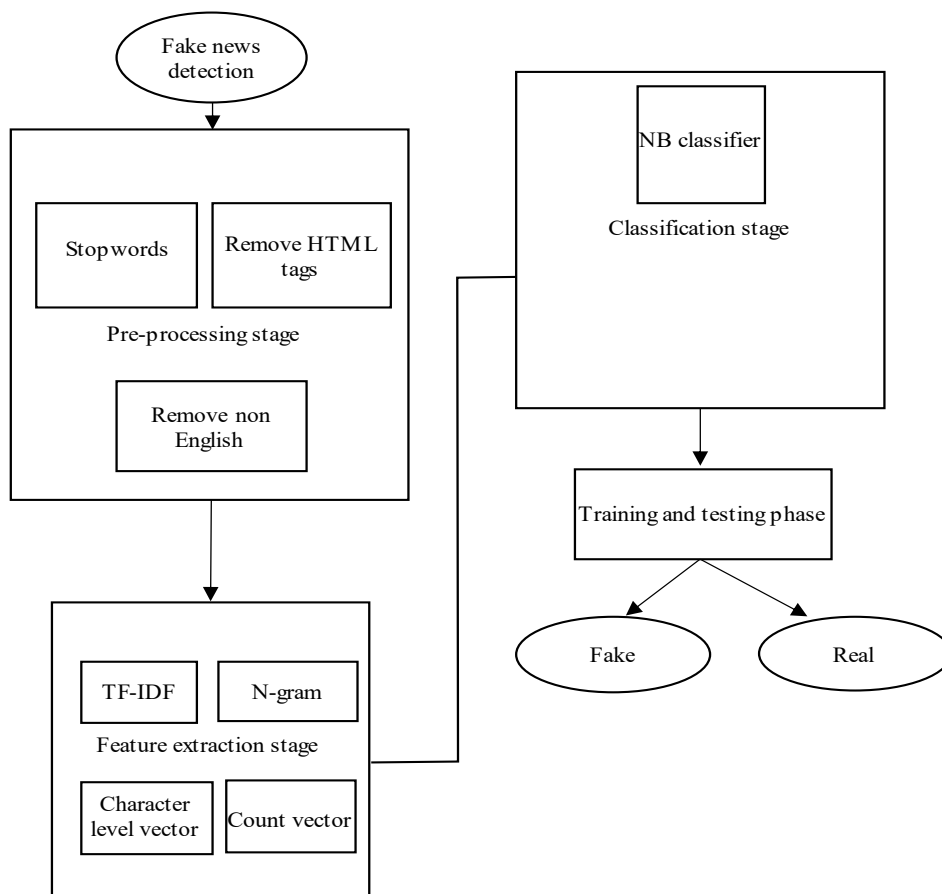


Figure 1. Fake news detection proposed model.

Feature Extraction Stage

Feature extraction entails converting text data into 0 and 1 vectors with new vectors generated from the sample text file. Various techniques exist for creating vectors, including the TF-IDF vectorizer, which calculates term frequency (TF) and inverse document frequency (IDF) [26, 27]. The tokenization vector breaks down texts into smaller segments, known as tokens, which are converted into vectors for machine comprehension. Tokenization can involve word tokenization, character tokenization, or sub-word tokenization.

TF-IDF vectorizer: The term frequency-inverse document frequency (TF-IDF) is a commonly used feature extraction method. This technique comprises two stages: the term frequency (TF) is initially calculated, followed by inverse document frequency (IDF) calculation [5].

$TF(t) = \text{Number of times term } t \text{ appears in a document} / \text{total number of terms in the document}$.

$IDF(t) = \text{Log}(\text{Total number of documents} / \text{Number of documents containing term } t)$

Classifier Stage

Several classifiers were employed in this study, including Naïve Bayes multinomial (NB), a classifier based on Bayes Naïve theory. NB is a straightforward supervised machine learning algorithm that categorizes objects into distinct classes [28]. It achieves high accuracy when handling large datasets, as shown in Figure 2.

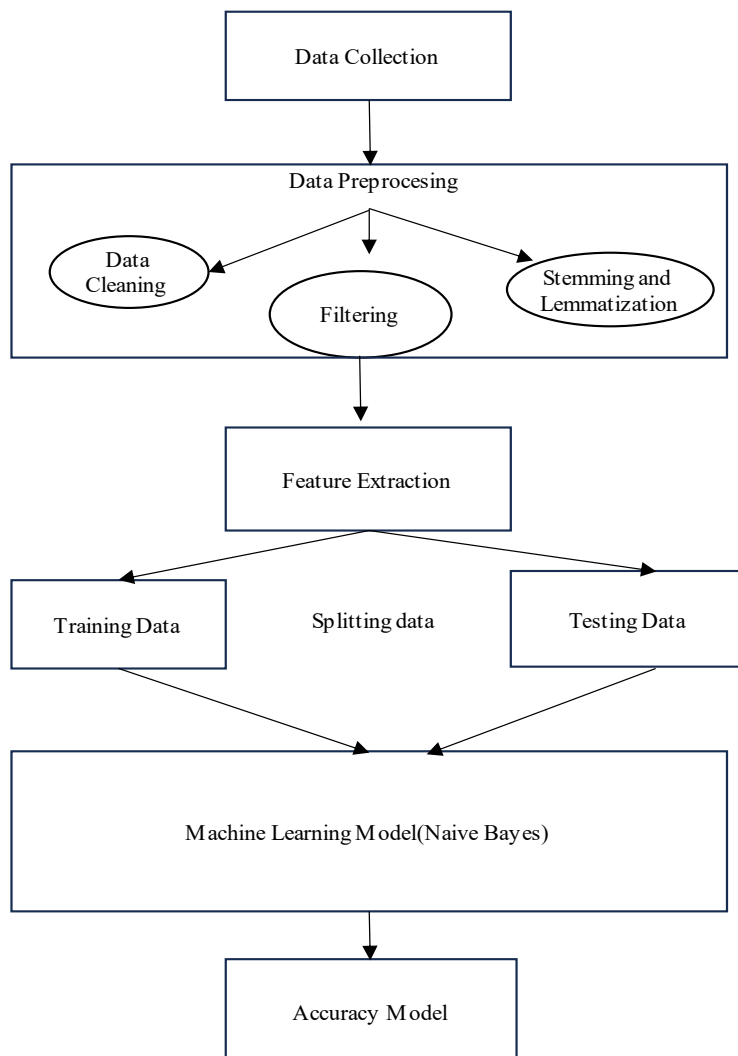


Figure 2. Data flow model.

IMPLEMENTATION

The practical implementation closely mirrors the framework and system design, delving into the intricacies of the system down to the code level. This section focuses on bringing earlier conceptualized ideas to fruition [29].

Data Collection

Online news is gathered from various sources including press agencies, search engines, and social media websites. For our project, we utilized a straightforward dataset comprising 6335 news articles classified as Fake or Real, which were subsequently stored in a CSV file. Dataset attributes include [30].

- *Id*: Unique identifier for the article
- *Title*: Headline of the article
- *Text*: Textual content of the article
- *Class Label*: Categorized as Fake or Real

Data Pre-processing

Given that data sourced from social media are often unstructured and informal, involving shortcuts, slang, and poor grammar, pre-processing is crucial for enhancing performance and reliability before utilizing it in predictive models [31].

Data Cleaning

Data can exist in either structured or unstructured formats. Structured data follow well-defined patterns, whereas unstructured data lacks a clear structure [32]. Cleaning involves:

1. *Punctuation removal*: Eliminating special characters that do not contribute to context.
2. *Tokenization*: Dividing text into units like words to impose structure.
3. *Removing stopwords*: Eliminating common, non-informative words.
4. *Stemming*: Reducing words to their base form to treat similar words uniformly.

Feature Generation

Various features, such as word count, repetition of distinctive words, and repetition of large words, were extracted from text data [33]. This is achieved by representing words in a way that captures their meaning and relationships, thereby facilitating the understanding of machine learning algorithms. Vectorization encodes text into a numerical form, typically through techniques such as the count vectorizer and TF-IDF.

Training the Model

Using machine learning classifiers, including Logistic Regression, Random Forest, Support Vector Machine, and Passive Aggressive, the model is trained after extracting features from preprocessed datasets [12]. The best-performing classifier, determined to be the Passive Aggressive classifier, was selected and stored for fake news classification. It takes the article input from the dataset and predicts news reliability.

- *True Positive (TP)*: Occurs when false information is correctly identified and classified as fake news.
- *True Negative (TN)*: Happens when true information is correctly identified and classified as true news.
- *False Negative (FN)*: Arises when true information is incorrectly classified as fake news.

Buntain et al. proposed a method for automating the detection of fake news on platforms such as Twitter by developing a deep understanding of the factors that contribute to the accuracy of credibility assessments in news reporting. This approach involves analyzing Twitter content using datasets from sources such as BuzzFeed's fake news dataset [34]. Through feature analysis, they identified the characteristics that are typically indicative of accuracy, drawing from both crowdsourced evaluations and editorial assessments.

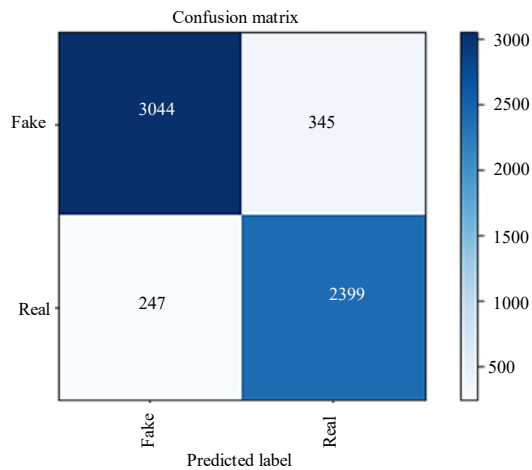


Figure 3. Confusion matrix.

FakeNewsDetection Home All user Prediction History Change Password Logout Welcome admin

Fake News Prediction Form

Write Your News Content

Trump Dxpessed Concern Over Anthony Weiner's illegal Access to Classified Info 2 Months age Truthfeed
 Stock market update: Nifty Realty index advances 0.19% in a weak markes

Predict

<p>FakeNewsDetection</p> <p>Enim minim veniam quis nostrud exercitation ullamco laboris nisl ut aliquo exed commodo consequat duis sed oute rure</p> <p>Phone contact@example.com Email: +54 305945234</p>	<p>The Company</p> <ul style="list-style-type: none"> > Site Map > Legal > Agent Admin > Careers > Affilote 	<p>International sites</p> <ul style="list-style-type: none"> > Amazon > Flipkart > Croma > Gadgets Now > Reliance Digital
---	--	---

FakeNewsDetection Home All user Prediction History Change Password Logout Welcome admin

Fake News Prediction Form

Witte Your News Content

Trump Expressed Concern Over Anthony Weiner's Illegal Access to Classified Info 2 Months age
 Truthreed Stuck market update: Nilty Realty index advances 0.19% in a weak market

Predict

Great! Given news is real.

<p>FakeNewsDetection</p> <p>Ervis minim veniam quis nostrud exercitation ullamco laboris nisi ut aliquip exea commodo consequat duis sed aute Irure.</p>	<p>The Company</p> <ul style="list-style-type: none"> > Site Map > Legal > Agent Admin 	<p>International sites</p> <ul style="list-style-type: none"> > Amazon > Flipkart > Flipkart
---	---	---

Figure 4. News detection.

Their methodology is built upon prior research that focuses on identifying highly retweeted conversation threads and extracting features from these threads to distinguish between credible and non-credible stories. However, it is important to note that this approach is limited in its applicability to only a subset of Twitter conversation threads, primarily those that receive significant levels of retweeting. Consequently, it may not be effective for the majority of less widely circulated tweets [35].

Confusion Matrix

A confusion matrix provides a structured representation of a classifier's performance on a given set of test data with known true values [36]. It offers a visual aid for gauging the effectiveness of the algorithm in making predictions. Presenting accurate and inaccurate predictions facilitates a detailed breakdown of each class. Essentially, a confusion matrix offers a concise summary of the prediction outcomes generated by the classification model [37]. It elucidates how the classifier might misinterpret certain data points, shedding light not only on its errors but also on the specific types of misclassifications occurring, as shown in Figure 3.

Model Evaluation

After the model was fitted, its performance was evaluated using a confusion matrix. True Positive (TP), True Negative (TN), False Negative (FN), and False Positive (FP) values were analyzed with the most accurate classifier chosen for news detection [38]. Precision, Recall, F1 Score, and accuracy metrics were used to evaluate the classifier performance, as shown in Figure 4.

CONCLUSION

The main goals of the Fake News Detection System project are multifaceted and aim to combat disinformation through a strategic combination of advanced machine learning and user-centered design. The first set of goals focuses on developing a robust machine learning model dedicated to fake news detection. This includes the implementation and training of the sophisticated multinomialNB algorithm, which is known for its advantages in text classification. In addition, the project aims to enhance the model's capabilities by exploring and integrating advanced natural language processing (NLP) techniques, ensuring increased accuracy and adaptability to the complex nuances of different language styles.

This study employed a multistep approach to detect fake news. Initially, XGBoost was utilized to prioritize variables based on their significance in fake news detection. Subsequently, a model was constructed using five classification algorithms (LR, NNET, RF, SVM, and CART), leveraging the factors identified through XGBoost. To enhance the performance and generalizability of each model, k-fold cross-validation was employed along with ablation studies to bolster robustness.

Machine learning, particularly supervised learning classification algorithms, has gained prominence in various research domains for predictive tasks. Ensemble Learning, a technique that combines multiple models, has become a powerful method for boosting algorithmic performance. XGBoost, a prominent ensemble model, often outperforms individual classification algorithms owing to its efficient parallel learning capabilities and regularization techniques, thus mitigating overfitting concerns.

The XGBoost model facilitates the identification of feature importance by considering the accuracy gains and frequencies of the variables across the entire decision tree. This provides insight into the variables' contributions to "important decisions" during the model-building process. By prioritizing the variables based on their significance, XGBoost mitigates noise and overfitting, thereby enhancing the accuracy and robustness of the fake news detection model.

Overall, this study aims to derive the key factors influencing fake news detection in tweet messages and construct an optimal model based on these insights. Through the application of XGBoost and ensemble learning techniques, this study seeks to improve the efficiency and accuracy of fake news identification, addressing the challenges posed by noisy data and overfitting in model construction.

REFERENCES

1. Wu C, Wu F, Huang Y, Xie X. Personalized news recommendation: Methods and challenges. *ACM Trans Inf Syst.* 2023;41:1-50. doi: 10.1145/3530257.
2. Su X, Sperli G, Moscato V, Picariello A, Esposito C, Choi C. An edge intelligence empowered recommender system enabling cultural heritage applications. *IEEE Trans Ind Informat.* 2019;15:4266-75. doi: 10.1109/TII.2019.2908056.
3. Zhou F, Xu X, Trajcevski G, Zhang K. A survey of information cascade analysis: Models, predictions, and recent advances. *ACM Comput Surv.* 2022;54:1-36. doi: 10.1145/3433000.
4. Gaillard S, Oláh ZA, Venmans S, Burke M. Countering the cognitive, linguistic, and psychological underpinnings behind susceptibility to fake news: A review of current literature with special focus on the role of age and digital literacy. *Front Commun.* 2021;6:661801. doi: 10.3389/fcomm.2021.661801.
5. Della Vedova ML, Tacchini E, Moret S, Ballarin G, DiPierro M, De Alfaro L. Automatic online fake news detection combining content and social signals. 2018 22nd Conference of Open Innovations Association (FRUCT), Jyväskylä, Finland, 2018, pp. 272-9. doi: 10.23919/FRUCT.2018.8468301.
6. Tandoc EC Jr, Lim ZW, Ling R. Defining “Fake News”: A typology of scholarly definitions. *Digit J.* 2018;6:137-53. DOI: 10.1080/21670811.2017.1360143.
7. Park M, Chai S. Constructing a user-centered fake news detection model by using classification algorithms in machine learning techniques. *IEEE Access.* 2023;11:71517-27. doi: 10.1109/ACCESS.2023.3294613.
8. Cybenko AK, Cybenko G. AI and fake news. *IEEE Intell Syst.* 2018;33:1-5. doi: 10.1109/MIS.2018.2877280.
9. Granik M, Mesyura V. Fake news detection using naive Bayes classifier. 2017 IEEE First Ukraine Conference on Electrical and Computer Engineering (UKRCON), Kyiv, Ukraine, 2017, pp. 900–903. DOI: 10.1109/UKRCON.2017.8100379.
10. Zhou S, Lin J, Tan L, Liu X. Condensed convolution neural network by attention over self-attention for stance detection in twitter. 2019 International Joint Conference on Neural Networks (IJCNN), Budapest, Hungary, 2019, pp. 1–8. doi: 10.1109/IJCNN.2019.8851965.
11. Taulé M, Martí MA, Rangel F, Rosso P, Bosco C, Patti V. Overview of the task on stance and gender detection in tweets on Catalan independence at IberEval 2017. In: *Proceedings of the Second Workshop on Evaluation of Human Language Technologies for Iberian Languages (IberEval 2017)*; 2017. p. 157-177. Available from: <https://ceur-ws.org/Vol-1881/Overview5.pdf>.
12. Weedon J, Nuland W, Stamos A. Information Operations and Facebook. Version 1.0. Facebook; 2017 Apr 27. Available from: https://i2.res.24o.it/pdf2010/Editrice/ILSOLE24ORE/ILSOLE24ORE/Online/_Oggetti_Embedded/Documenti/2017/04/28/facebook-and-information-operations-v1.pdf.
13. Wei P, Mao W, Zeng D. A target-guided neural memory model for stance detection in twitter. 2018 International Joint Conference on Neural Networks (IJCNN), Rio de Janeiro, Brazil. 2018. pp. 1–8. DOI: 10.1109/IJCNN.2018.8489665.
14. Helmstetter S, Paulheim H. Weakly supervised learning for fake news detection on Twitter. 018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), Barcelona, Spain, 2018. pp. 274-277. DOI: 10.1109/ASONAM.2018.8508520.
15. Mandical RR, Mamatha N, Shivakumar N, Monica R, Krishna AN. Identification of fake news using machine learning. 2020 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT), Bangalore, India, 2020, pp. 1-6, doi: 10.1109/CONECCT50063.2020.9198610.
16. Augenstein I, Rocktäschel T, Vlachos A, Bontcheva K. Stance Detection with Bidirectional Conditional Encoding. [Preprint]. arXiv:1606.05464. 2016 Jun 17. DOI: <https://doi.org/10.48550/arXiv.1606.05464>.
17. Ahmad I, Yousaf M, Yousaf S, Ahmad MO. Fake news detection using machine learning ensemble methods. *Complexity.* 2020;2020:1–11. doi: 10.1155/2020/8885861.

18. Kaur S, Kumar P, Kumaraguru P. Automating fake news detection system using multi-level voting model. *Soft Comput.* 2020;24:9049-69. doi: 10.1007/s00500-019-04436-y.
19. Ghosh S, Desarkar MS. Class specific TF-IDF boosting for short-text classification: Application to short-texts generated during disasters. *Companion Proceedings of the Web Conference 2018; 2018 Apr 23-27; Lyon, France. New York: International World Wide Web Conferences Steering Committee; 2018.* p. 1629-37. doi: 10.1145/3184558.3191621.
20. Elfardy H, Diab M. CU-GWU perspective at SemEval-2016 task 6: Ideological stance detection in informal text. In: Bethard S, Carpuat M, Cer D, Jurgens D, Nakov P, Zesch T, editors. *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016); 2016 Jun; San Diego, California. Association for Computational Linguistics; 2016.* p. 434–9. doi: 10.18653/v1/S16-1070.
21. Laboreiro G, Sarmento L, Teixeira J, Oliveira E. Tokenizing micro-blogging messages using a text classification approach. In: *Proceedings of the Fourth Workshop on Analytics for Noisy Unstructured Text Data; 2010 Jul 11-16; Toronto, ON, Canada. New York: Association for Computing Machinery; 2010.* p. 81-8. doi: 10.1145/1871840.1871853.
22. Aldwairi M, Alwahedi A. Detecting fake news in social media networks. *Procedia Comput Sci.* 2018;141:215–22. doi: 10.1016/j.procs.2018.10.171.
23. Masood AD, Abdulazeez AM, Zeebaree DQ. Machine learning supervised algorithms of gene selection: A review. *Mach Learn.* 2020;62(3):233–244.
24. Patra BG, Das D, Bandyopadhyay S. JU_NLP at SemEval-2016 Task 6: Detecting stance in tweets using support vector machines. In: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016); 2016 Jun; San Diego, California. San Diego: Association for Computational Linguistics; 2016.* p. 440-444. DOI: 10.18653/v1/S16-1071. Available from: <https://aclanthology.org/S16-1071/>
25. Buntain C, Golbeck J. Automatically identifying fake news in popular twitter threads. *2017 IEEE International Conference on Smart Cloud (SmartCloud), New York, NY, USA. 2017.* pp. 208-215. doi: 10.1109/SmartCloud.2017.40.
26. Dong L, Wesseloo J, Potvin Y, Li X. Discrimination of mine seismic events and blasts using the fisher classifier, Naive Bayesian classifier and logistic regression. *Rock Mech Rock Eng.* 2016;49:183-211. doi: 10.1007/s00603-015-0733-y.
27. Mohammad S, Kiritchenko S, Sobhani P, Zhu X, Cherry C. SemEval-2016 Task 6: Detecting stance in tweets. In: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016); 2016 Jun; San Diego, California. San Diego: Association for Computational Linguistics; 2016.* p. 31-41. DOI: 10.18653/v1/S16-1003. Available from: <https://aclanthology.org/S16-1003/>
28. Vereshchaka A, Cosimini S, Dong W. Analyzing and distinguishing fake and real news to mitigate the problem of disinformation. *Comput Math Organ Theory.* 2020;26:350-64. doi: 10.1007/s10588-020-09307-8.
29. Al-Ghadir AI, Azmi AM, Hussain A. A novel approach to stance detection in social media tweets by fusing ranked lists and sentiments. *Inf Fusion.* 2021;67:29-40. doi: 10.1016/j.inffus.2020.10.003.
30. Conroy NK, Rubin VL, Chen Y. Automatic deception detection: Methods for finding fake news. *Proc Assoc Inf Sci Technol.* 2015;52:1–4. doi: 10.1002/pra2.2015.145052010082.
31. Somasundaran S, Wiebe J. Recognizing stances in ideological on-line debates. In: Inkpen D, Strapparava C, editors. *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text; 2010 Jun 1-2; Los Angeles, CA. Stroudsburg (PA): Association for Computational Linguistics; 2010.* p. 116-24. Available from: <https://aclanthology.org/W10-0214/>
32. Konjengbam A, Ghosh S, Kumar N, Singh M. Debate stance classification using word embeddings. In: Ordonez C, Bellatreche L, editors. *Big Data Analytics and Knowledge Discovery. DaWaK 2018. Lecture Notes in Computer Science, vol 11031. Cham: Springer; 2018.* p. 377-91. DOI: https://doi.org/10.1007/978-3-319-98539-8_29.

-
33. Faulkner A. Automated classification of stance in student essays: An approach using stance target information and the Wikipedia link-based measure. In: Proceedings of the Twenty-Seventh International Florida Artificial Intelligence Research Society Conference (FLAIRS 2014); 2014; Palm Beach, Florida. Association for the Advancement of Artificial Intelligence; 2014. p. 174–9.
 34. Vlachos A, Riedel S. Identification and verification of simple claims about statistical properties. *Rt-PA 2015 Conf Empir Methods Nat Lang Process. Assoc Comput Linguist.* p. 2596–601. doi: 10.18653/v1/D15-1312.
 35. Harabagiu S, Hickl A, Lacatusu F. Negation, contrast and contradiction in text processing. In: Proceedings of the AAAI Conference on Artificial Intelligence. Boston, MA. Technical Papers. 2006. p. 755-762.
 36. Sen A, Sinha M, Mannarswamy S, Roy S. Stance classification of multi-perspective consumer health information. In: Proceedings of the ACM India Joint International Conference on Data Science and Management of Data; 2018; p. 273–81. DOI: 10.1145/3152494.3152518.
 37. Silverman C. Lies, damn lies, and viral content. *Columbia Journalism Rev.* 2015. DOI: 10.7916/D8Q81RHH. Available from: https://www.cjr.org/tow_center_reports/craig_silverman_lies_damn_lies_viral_content.php
 38. Vychezhnanin SV, Kotelnikov EV. Stance detection based on ensembles of classifiers. *Program Comput Softw.* 2019;45:228–40. doi: 10.1134/S0361768819050074.